Are carbon emissions associated with stock returns?

Rejoinder to March 2023 comment by Patrick Bolton and Marcin Kacperczyk

Jitendra Aswani
Harvard University

Aneesh Raghunandan
London School of Economics

Shiva Rajgopal
Columbia University

April 2023

Introduction and acknowledgments

Before directly addressing the *Response to critique and further elaboration* (hereafter "Comment"), written by Patrick Bolton and Marcin Kacperczyk (BK), a few thanks are in order. First, we thank the editor, Alex Edmans, both for providing BK the opportunity to directly comment on our paper as well as giving us the opportunity to respond. We believe that open intellectual debate is vital especially given the topic at hand. Directly facilitating such a conversation in conjunction with the paper's publication in *RF* seems to us a good way to showcase various perspectives with equal prominence. Second, we thank BK for taking up the offer to engage with us within this context. We believe direct engagement is the best way to advance scholarly debate on whether and how investors consider carbon emissions information.

In what follows, we write our rejoinder to BK's Comment in the format of a response memo, i.e., we reproduce BK's main points (with some paraphrasing) and respond to each individually in turn. We believe this is the clearest – and, importantly, most parsimonious – way to provide our thoughts.

Response to BK's *Comment*

Point #1: Unscaled emissions is in fact the 'correct' measure of carbon risk

We agree with BK that this is a matter of how to interpret empirical results, in the sense that both their studies and ours consistently find a relation between returns and unscaled emissions, but not between returns and emissions intensity. We, however, disagree with BK's argument that unscaled emissions are the best way to assess carbon transition risk *for individual firms* (we do agree that unscaled emissions are the appropriate way to assess society as a whole). BK provide three arguments in their *Comment*; we discuss each of these in turn.

First, BK argue that relying on carbon intensity means that some large firms may be seem as more environmentally friendly than smaller firms despite contributing a greater share of the overall economy's emissions. We do not see this as a problem, under the assumption that it is not necessarily desirable to force firms to reduce their levels of *production*. Assuming consumers demand a specific quantity of a good, it is overall 'greener' for the economy if firms with lower emissions intensity produce more of the good.

For example, if consumers demand 100 units of a good, two firms producing 50 units each and emitting 2 metric tons of CO_2 per unit result in less overall pollution than ten firms producing 10 units each and emitting 3 metric tons of CO_2 per unit. Yet, according to BK's argument, the latter case ought to result in lower carbon risk because each of the ten firms has lower total emissions (30 tons) than the two firms in the former case (100 tons) – *even though overall emissions in the economy are 50% higher*.[1] Thus, unless BK's intention is to argue that overall production in the economy must decrease, we do not think total emissions are an appropriate measure of *firm*-level carbon risk.

We find the interpretation of the FT rankings of Europe's Climate Leaders provided in BK's *Comment* to be puzzling; it seems to us that firms that grow their sales and production faster than their emissions ought to be celebrated as reducing their carbon risk, not chastised.

Second, BK provide the example of Net Zero pledges. With respect to such pledges, we note that when total emissions are zero, so too is emissions intensity. Thus, a Net Zero pledge can just as easily be interpreted as a pledge to ultimately cut emissions *intensity* to zero. Indeed, many of the interim targets disclosed by firms on the path to net zero are stated in terms of emissions intensity. Moreover, Net Zero pledges themselves do not necessarily imply a firm's intention to cut actual emissions from operations and the end use of products (i.e., the sum of scopes 1, 2, and 3) to zero; many firms approach such pledges by purchasing carbon credits to offset existing emissions.

Third, BK state that dividing emissions by sales revenue is likely to induce some noise. We acknowledge this point. We argue, however, that this is less problematic than including both total emissions and a proxy for size as two separate variables on the right-hand side of a regression; the latter approach introduces severe multicollinearity. In an economic sense, this multicollinearity arises because total emissions and total output are both derived from the same set of raw inputs and, hence, an increase in one is mechanically highly correlated with an increase in the other. We discuss this issue further in our response to Point #3 below. More

[1] Moreover, as we highlight in our paper, such reasoning leads to a logical inconsistency. Consider Firm A that produces \$100 worth of goods and emits 10 metric tons of CO_2, and Firm B that produces \$50 worth of goods and emits 7 metric tons of CO_2. BK argue that Firm B ought to be thought of as 'greener'. However, if Firm A split itself into two equally-sized parts, each would then produce \$50 worth of goods and emit 5 metric tons of CO_2 – making *both* new firms 'greener' than Firm B despite the total amount of emissions in the economy not changing!

generally, claiming that any variable highly correlated with size, such as emissions, is priced in equity returns independent of size, is a tenuous proposition given the long history of robust correlations between size and returns.

As a final note, we believe we do not 'overlook' the evidence linking unscaled emissions to stock returns as BK's *Comment* claims. Indeed, our own study first *directly replicates* these analyses before highlighting potential issues therein.

Point #2: Stock returns are in fact correlated with emissions disclosed by firms

BK note that their own studies *do* document a correlation between carbon emissions and stock returns when considering only firm-disclosed (rather than vendor-estimated) emissions. We acknowledge this point, both here and in our published paper (see Section 5.2). Indeed, as we were revising our own study, we initially found this discrepancy puzzling, and so we investigated in some detail.

The results of this investigation are discussed in Section 5.2 of the published paper and the corresponding Online Appendix Table OA3. We also briefly summarize here. More specifically, in these analyses we varied (i) the way in which we constructed the disclosure measure and (ii) the choice of industry fixed effect used in our regressions.

With respect to (i), Trucost provides (as of this writing) 29 different values for its 'data source' variable. We label these 29 values as disclosed or estimated in two different ways, one of which is significantly more conservative than the other; these are detailed in Sections 4 and 5.2, and corresponding Footnotes 9 and 13 in the manuscript.

With respect to (ii), we consider four types of fixed effects: three of which are standard in academic research (GICS, SIC, Fama-French) as well as a fourth (based on Trucost industry definitions) which, while uncommon in academic work, mimics the BK studies alluded to in the *Comment*.

Our main finding in this regard is that BK's result that US disclosed emissions are correlated with stock returns, appears to be *extremely* sensitive to the way in which the disclosure measure is constructed and fixed effects are chosen. As Table OA3 shows, we only find a positive and significant effect when we precisely mimic BK's specification choice (conservative disclosure measure, Trucost industry fixed effects). No other reasonable variation in econometric specification yields evidence of a relation between disclosed emissions and stock returns. To that end, while we do not dispute the precise result that BK show across their studies, we are concerned about its sensitivity to specification choices and, hence, its generalizability.

Point #3: Variations in emissions are not driven by firm fundamentals

In their third comment, BK first state that they do control for firm size. As we highlight below in Table R1, however, *controlling* for size is not necessarily an alternative to *scaling* by size. Given the high univariate correlation between emissions and firm size, the first-order effect of *controlling* for size in a linear regression of returns on total emissions is to introduce significant multicollinearity, which in turn increases the risk of spurious inferences.

As evidence of such multicollinearity, in Table R1 below we reproduce columns (4) – (9) of Table 5 of our paper, where we regress stock returns on unscaled carbon emissions. Now, however, we provide an additional test statistic at the bottom of each column: the variance inflation factor (VIF), commonly used to assess the issue of multicollinearity within a regression, for the coefficient on the respective log emissions variables.

By default, the VIF when including only emissions on the right-hand side is 1. However, including size alone – to say nothing of other control variables – leads to VIFs of 5.24, 4.59, and 5.63 with just the addition of size as a control variable. This is consistent with the high univariate correlations between size and unscaled emissions documented in the manuscript. While there is no bright-line theoretical VIF threshold for multicollinearity, a general rule of thumb is that a VIF above 4 indicates significant multicollinearity. The numbers provided above all far exceed this threshold and may well explain why BK only find a relation between emissions and returns when also controlling for size (but do not find a univariate relation). To this end, controlling for size in a regression of stock returns on unscaled emissions may create more problems than it solves.

As part of their third comment, BK also allude to links between emissions and both market-to-book and P/E ratios. However, given the high univariate correlation between size and emissions it is not clear that this captures anything beyond potential relations between size and both P/E ratios and market-to-book.

Other/minor point: Paris Agreement analyses

In their *Comment* BK state that they find the relation between emissions and unscaled returns is stronger after the 2015 Paris agreement relative to before. This appears to refer to Table 14, discussed within Section 3.6, of their published JFE paper about the US setting. We note, however, that they also state the following in that paper (bolding ours):

"However, an important caveat is that ***our sample increases after 2015****, so that the difference in returns pre and post Paris could be attributed to the new firms that were added to our sample. We explore this possibility in the Online Appendix and indeed find that the increase in return premium is mostly due to the addition of the new firms. The results are reported in Table A.6. We find that* ***when we exclude the new firms, the carbon premium becomes insignificant in the two years following the Paris agreement****."*

As we discuss in our own paper, in 2015 Trucost substantially expanded the set of firms for which it provided emissions data; this is almost certainly the reason for the post-2015 sample increase that BK allude to in the above text. However, nearly all the newly covered firms only had *estimated* emissions figures. Thus, evidence of an increased 'carbon premium' post-2015 may simply reflect a change in the sample proportion of firms with estimated emissions figures rather than an actual economic effect. BK's own analyses (in their Table A.6) appear to support this conclusion.

Table R1: Stock Returns and Carbon Emissions

This table provides results from regressions of monthly stock returns on emissions and reflects columns (4) – (9) of Table 5 in our published paper. For consistency with our paper, we number columns (4) – (9) rather than (1) – (6). In columns (4) – (6) we estimate regressions without controls but with industry fixed effects; in columns (7) – (9) we add firm size. We report standard errors in parentheses beneath coefficient estimates. In all panels, *, **, and *** indicate statistical significance at 10% ,5%, and 1% levels, respectively. We report a new test statistic in this table relative to our published paper: the Variance Inflation Factor for the coefficient on either log scope 1, 2, or 3 emissions, labelled as "VIF for emissions" in the table below.

	(4)	(5)	(6)	(7)	(8)	(9)
VARIABLES	Ret	Ret	Ret	Ret	Ret	Ret
Log Scope 1	-0.048			0.089**		
	(0.039)			(0.040)		
Log Scope 2		-0.047			0.139***	
		(0.041)			(0.049)	
Log Scope 3			-0.044			0.245***
			(0.051)			(0.069)
Firm Size				-0.238***	-0.292***	-0.386***
				(0.087)	(0.097)	(0.106)
Observations	178,354	178,354	178,354	178,354	178,354	178,354
R^2	0.190	0.190	0.190	0.191	0.191	0.191
VIF for emissions	1.00	1.00	1.00	5.24	4.59	5.63
Other Controls	No	No	No	No	No	No
Industry	Yes	Yes	Yes	Yes	Yes	Yes
Month-Year	Yes	Yes	Yes	Yes	Yes	Yes